Learning to Compute the Plane of Symmetry for Human Faces

Jia Wu[1], Raymond Tse[2], Carrie L. Heike[1,2], Linda G. Shapiro[1]
1. University of Washington, Seattle, WA, 98195, U.S.A.
2. Seattle Children's Hospital Craniofacial Center, Seattle, WA, 98105, U.S.A.
{jiawu, shapiro}@cs.washington.edu
{raymond.tse, carrie.heike}@seattlechildrens.org

ABSTRACT

Facial symmetry analysis is complex in both computer vision and medicine. This paper presents a method to compute the plane of symmetry for 3D meshes of the human head and face through learning. The two steps of processing include: 1) landmark-related region detection and 2) symmetry plane computation in the learning stage, which uses the landmarks and the standard symmetry planes identified by medical experts for training. Experimental results show that our method performs better than the existing mirror method [1], and is robust to rotation and noise.

Categories and Subject Descriptors

I.2.6 [**Computing Methodologies**]: Artificial Intelligence—*Learning*; I.4.9 [**Computing Methodologies**]: Image Processing and Computer Vision—*Applications*; J.3 [**Computer Applications**]: Life and Medical Sciences

Keywords

3D mesh, facial symmetry, learning

1. INTRODUCTION

Symmetry is a common phenomenon in nature, and there are many applications based on analysis of symmetry in computer vision and graphics, such as image editing [8, 12], symmetrizing deformations for 3D data [13], and data repair [15, 19]. Researchers from computer vision and medical imaging share an interest in computation of human face symmetry. Symmetry analyses have been used for studying facial attractiveness [10], quantification of degree of asymmetry in individuals with craniofacial birth defects (before and after corrective surgery) [17], and analysis of facial expression for human identification [11]. Benz *et al.* [1] introduced a common method for the 3D facial symmetry analysis in which the original data is mirrored at an arbitrary plane, then the original and the mirrored mesh are registered using the

iterated-closest-point algorithm [2]. Based on the registered data, the symmetry plane is determined from the centers of associated points. This method is reliable when the data is properly aligned [7], but the results rely heavily on the choice of the initial plane about which the data is mirrored and the method is not robust to noise [1].

In this paper, a two step approach for learning the plane of symmetry is proposed. The first step is point-based region detection, while the second step uses appropriate regions to determine symmetry properties.

2. METHOD

The proposed method is based on machine learning. There are two stages of processing: landmark-related region detection and symmetry plane computation.

2.1 Data Format and Training Set

The 3D data used in our study were collected by investigators in the Craniofacial Center of Seattle Children's Hospital (SCH) using the 3dMDcranial[TM] imaging system. The data are 3D meshes without facial texture maps. 40 head meshes from 40 individuals were used for training. Two medical experts labeled a set of landmarks on the original textured 3D models and the average value for each landmark was used for training. All training data were aligned to face forward by a semi-automatic method [18].

2.2 Learning Landmark-related Points

The regions used in this work are related to facial landmarks. The landmarks used here are a subset of the landmarks used for craniofacial anthropometric measurements for clinical and research applications [9]. Seven midline points and four paired landmarks are used, as illustrated in Fig. 1 (a). Multiple points in a small region around a certain landmark are used for training instead of a single point, because they have similar properties and are more reliable. In the training set, points around expert-marked landmark k are labeled as positive samples for that landmark, and points that are not near landmark k form the pool of negative samples for that landmark.

Local properties are calculated to form a signature vector at each surface point. The mean curvature H and Gaussian curvature K [3] were estimated on both the original data and on a sequence of smoothed data. Hamming windows [6] with multiple sizes were used to smooth the data. Signature vectors were used to train a REPTree (the Reduced Error Pruned Tree) [16] classifier M_L to learn the interesting points on the 3D surface mesh. The WEKA im-
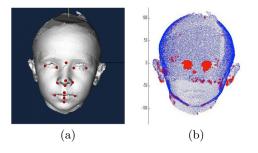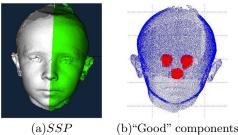
Figure 1: (a) The 3D head mesh data with landmarks labeled by medical experts (b)Interesting region prediction for inner corner of eye

plementation of REPTree was used in all experiments [5]. After training was complete, each model $M_{L,k}$ was able to label every point on a 3D head as either interesting or not for each landmark k and to provide a confidence score for that landmark. A threshold, $T = 0.98$, was applied to the confidence scores. Figure 1(b) shows sample results of landmark classification for inner corner of eye (en). The false positives will be removed in the next stage.

2.3 Learning Regions to Use in Computing the Location of Symmetry Plane

The second training step, which is used to build the symmetry model, is based on the interesting regions. The points obtained from the last step are extracted and used to form connected components. Components less than 10 points are removed. For each component m in landmark model k, the center coordinates C_m are extracted, as well as the eigenvalues λ_{m1}, λ_{m2} and λ_{m3} of the matrix formed by the 3D coordinates of all the points in component m. Eigenvalues are used because they reflect the size of the interesting regions and are invariant to rotation.

In the training set, the standard symmetry plane SSP is determined from the 11 ground-truth landmarks. This plane is selected by minimizing the distance of the plane to each of the single landmarks and to the average of each of the paired landmarks as shown in Figure 2(a).

(a)SSP   **(b)"Good" components**

Figure 2: Training symmetry model.(a) Standard symmetry plane (SSP) ground truth (b) a "good" single component and a good pair of components

A single component which lies on the symmetry plane and pairs of components that are approximately symmetric to the symmetry plane are marked as positive samples, which means they are suitable for use in symmetry plane construction. The rest of the components are negative sam-

ples. Figure 2(b) shows a good single component and a good pair of components.

The feature vector to describe single component m includes the number of points in component m, as well as the 3 eigenvalues. The feature vector for paired components m and n is designed to capture the difference between these two components. It contains the distance between their areas, eigenvalues, and centers. Using the features $Csingle$ along with positive and negative samples for $Goodsingle$, a REPTree classifier M_S was trained to learn if a component is suitable for construction of the plane of symmetry. Similarly, using the features $Cpair$ and positive and negative samples for $Goodpair$, another REPTree classifier M_P was trained to learn if a pair of components is suitable.

The above two steps are performed on the training set. Now for each landmark k, there is one landmark model $M_{L,k}$, one model for determining good single components $M_{S,k}$ and one model for determining good pairs of components $M_{P,k}$.
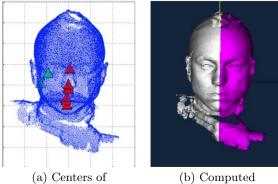
2.4 Computing the Symmetry Plane

For new data classification, the steps to identify "good" components are:

1. Smooth the data by Hamming windows of the designed sizes

2. Compute curvatures for original data and all the smoothed data for each point on the 3D head mesh

3. For every landmark type k

   Identify interesting regions: use landmark model $M_{L,k}$ to get a landmark-related confidence score for every point on the mesh.

   Compute features for components: pick all points with score larger than $T = 0.98$ and form components of 10 or more points. Compute $Csingle$ and $Cpair$ for all components.

   Classify components: run the model for good single components $M_{S,k}$ and the model for good pair components $M_{P,k}$ to label components as good singles, good pairs, or neither.

4. Accumulate the good singles and good pairs for all the landmark models

After all the "good" components have been classified, a symmetry plane is determined that satisfies two criteria: 1) the plane should be perpendicular to the line connecting each good pair and lie along the centers of good pairs. 2), the plane should be close to the centers of good single components. The plane P was chosen to minimize the function f

$$f = \sum D(P, C_s) + \sum D(P, C_p) + \gamma Angle(P, C_p)$$

where $D(P, C_s)$ is the distance of plane P to the centers of good single components, $D(P, C_p)$ is the distance of plane P to the centers of good pairs of components, $Angle(P, C_p)$ stands for the angle between the normal of plane P and the vector connecting a good pair of components, and γ is a weight. In order to obtain the optimal plane, the RANSAC [4] algorithm was applied to get rid of false positive "good" components, which are not consistent with the others. (See Figure 3.)

(a) Centers of
"Good" components

(b) Computed
symmetry plane

Figure 3: Centers of good singles and good pairs and the plane of symmetry (a) RANSAC prunes out false positives (green) (b) the plane of symmetry is shown on an original 3D mesh

3. EXPERIMENTS AND RESULTS

Several experiments on different 3D data sets including heads at different orientations, facial images with significant asymmetry, and heads with asymmetrical shape were run. We compared our method with the commonly used mirror approach [1, 14].

3.1 Performance on Data Set one

This data set consists of 15 3D heads and landmarks labeled by medical experts. Using the landmarks, the standard symmetry plane (SSP) was obtained for each 3D head and served as its ground truth.

There are two metrics used to evaluate the result of the computed symmetry plane (CSP), both from our learning method and the mirror method. The first metric is the angle Dif_a between the SSP and the CSP. The second metric F-measure reflects how the CSP and the SSP agree with each other on a point by point basis. The closer the F-measure value is to 1, the closer are the CSP and the SSP. Figure 4 gives the comparison of our learning method and the mirror method on this data set. Our method does not result in large errors and outperforms the mirror method in 60% of the cases in this data set. The mirror method shows very poor performance in two cases, both of which are caused by improper initial position of the data.

3.2 Performance on Noisy, Rotated Data Set

With different sizes of volume blocks placed on one or both sides of the face, one individual's 3D image was taken ten times to form the second data set. The landmarks labeled by medical experts were used to compute the SSP (ground truth) for each version of the data. In order to test each method's capability to handle random rotations, all data were rotated in the yaw and roll directions for $30°$, $60°$ and $90°$. As such, 16 orientations were created per head in the original data set, leading to 160 head meshes.

Our learning method and the mirror method were applied to the rotated data. Figure 5 shows the capability for both methods to handle rotated heads. The horizontal axis is the rotation angle label. The vertical axes represent the angle between the CSP and the SSP in (a) and F-measure in
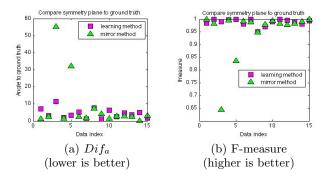

(a) Dif_a
(lower is better)

(b) F-measure
(higher is better)

Figure 4: Comparison of our learning approach to the mirror approach (a) Angle differences (b) F-measure for point differences

(a) Dif_a
(lower is better)

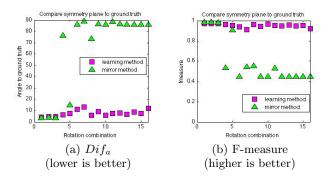(b) F-measure
(higher is better)

Figure 5: Comparison of our learning approach to previous used mirror approach according to different rotation angles (a) Angle differences (b) F-measure

(b). The numbers are the average of the 10 heads at each particular rotation. The mirror method performs best with yaw or roll in one direction within $30°$. The learning method produced good results despite the effect of noise (asymmetry created using volume blocks on cheeks) and rotation.

3.3 Performance on Cleft Data Set

The most challenging data set in our study included 3 individuals with unrepaired clefts of the lip and palate, 19 individuals who had undergone lip repair, and 18 individuals with other craniofacial conditions, such as skull asymmetry. The first column of Fig. 6 gives an example of a child with a bilateral cleft of the lip and palate. Our learning method can find the symmetry plane across the center of the mouth even though it has an abnormal shape, while the mirror method does not find the center well. In the second column of Fig. 6, the 3D mesh is missing data points on the lower right. Our method can correctly find the symmetry plane. In the last column of Fig. 6, the individual has an asymmetrical head shape at the back of his head (posterior plagiocephaly), which adversely affects the result of the mirror method, but our learning method is able to perform correctly despite this asymmetry. While there is no ground truth SSP available for this data set, medical experts noted that our results were highly correlated with their clinical impressions of the midline plane of symmetry.
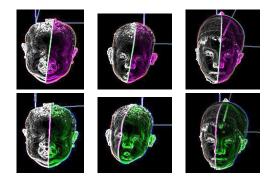
Figure 6: sample results of learning method (top row) and mirror method (bottom row) of a bi-lateral cleft, missing data and asymmetrical head shape.

4. CONCLUSIONS

We have developed a new methodology for automatically computing the plane of symmetry for human faces. Our method uses two stages of learning: landmark-related region learning and symmetry region learning. Our experimental results show that it performs better than the commonly used mirror method. Furthermore, our method does not require pose-alignment processing and is invariant to rotation.

Future work includes quantification of facial asymmetry, which would have applications for medical research, and could be used for more general tasks, such as face recognition and facial expression analysis. In addition, this learning framework could be transferred to other 3D models to compute the plane of symmetry or to look for partial or global symmetry in 3D data.

5. ACKNOWLEDGMENTS

6. REFERENCES

[1] M. Benz, X. Laboureux, T. Maier, E. Nkenke, S. Seeger, F. Neukam, and G. Häusler. The symmetry of faces. In *Proceedings of Vision, Modeling, and Visualization, G. Girod, H. Niemann, T. Ertl, B. Girod, and H.-P. Seidel, eds.(Akademische Verlagsgesellschaft, 2002)*, pages 43–50, 2002.

[2] P. Besl and N. McKay. A method for registration of 3-D shapes. *IEEE Transactions on pattern analysis and machine intelligence*, pages 239–256, 1992.

[3] K. Chang, W. Bowyer, and P. Flynn. Multiple nose region matching for 3D face recognition under varying facial expression. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 28(10):1695–1700, 2006.

[4] M. Fischler and R. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981.

[5] E. Frank, M. Hall, L. Trigg, G. Holmes, and I. Witten. Data mining in bioinformatics using Weka. *Bioinformatics*, 20(15):2479, 2004.

[6] F. Harris. On the use of windows for harmonic analysis with the discrete Fourier transform. *Proceedings of the IEEE*, 66(1):51–83, 1978.

[7] J. Hartmann, P. Meyer-Marcotty, M. Benz, G. Häusler, and A. Stellzig-Eisenhauer. Reliability of a method for computing facial symmetry plane and degree of asymmetry based on 3D-data. *Journal of Orofacial Orthopedics/Fortschritte der Kieferorthopädie*, 68(6):477–490, 2007.

[8] K. Huang, W. Hong, and Y. Ma. Symmetry-based photo-editing. *Pattern Recognition*, 38(6):825–834, 2005.

[9] J. Kolar and E. Salter. *Craniofacial anthropometry: Practical measurement of the head and face for clinical, surgical, and research use*. CC Thomas, 1997.

[10] M. Komori, S. Kawamura, and S. Ishihara. Averageness or symmetry: Which is more important for facial attractiveness? *Acta psychologica*, 131(2):136–142, 2009.

[11] Y. Liu, K. Schmidt, J. Cohn, and S. Mitra. Facial asymmetry quantification for expression invariant human identification. *Computer Vision and Image Understanding*, 91(1-2):138–159, 2003.

[12] G. Loy and J. Eklundh. Detecting symmetry and symmetric constellations of features. *Computer Vision–ECCV 2006*, pages 508–521, 2006.

[13] N. Mitra, L. Guibas, and M. Pauly. Symmetrization. In *ACM SIGGRAPH 2007 papers*, page 63. ACM, 2007.

[14] M.-M. P, A. GW, G. AB, and S.-E. A. Impact of facial asymmetry in visual perception: a 3-dimensional data analysis. *American journal of orthodontics and dentofacial orthopedics*, pages 137(2):168.e1–8; discussion 168–9, 2010.

[15] M. Pauly, N. Mitra, J. Wallner, H. Pottmann, and L. Guibas. Discovering structural regularity in 3D geometry. *ACM Transactions on Graphics-TOG*, 27(3):43–43, 2008.

[16] J. Quinlan. *C4. 5: programs for machine learning*. Morgan Kaufmann, 1993.

[17] I. Stauber, E. Vairaktaris, A. Holst, M. Schuster, U. Hirschfelder, F. Neukam, and E. Nkenke. Three-dimensional analysis of facial symmetry in cleft lip and palate patients using optical surface data. *Journal of Orofacial Orthopedics/Fortschritte der Kieferorthopädie*, 69(4):268–282, 2008.

[18] K. Wilamowska, L. Shapiro, and C. Heike. Classification of 3D face shape in 22q11. 2 deletion syndrome. In *Biomedical Imaging: From Nano to Macro, 2009. ISBI'09. IEEE International Symposium on*, pages 534–537. IEEE, 2009.

[19] K. Xu, H. Zhang, A. Tagliasacchi, L. Liu, G. Li, M. Meng, and Y. Xiong. Partial intrinsic reflectional symmetry of 3d shapes. *ACM Trans. Graph*, 28(5):1–10, 2009.